Incremental Control Synthesis in Probabilistic Environments with Temporal Logic Constraints

Alphan Ulusoy[*]   Tichakorn Wongpiromsarn[†]   Calin Belta[*]

Abstract— In this paper, we present a method for optimal control synthesis of a plant that interacts with a set of agents in a graph-like environment. The control specification is given as a temporal logic statement about some properties that hold at the vertices of the environment. The plant is assumed to be deterministic, while the agents are probabilistic Markov models. The goal is to control the plant such that the probability of satisfying a syntactically co-safe Linear Temporal Logic formula is maximized. We propose a computationally efficient incremental approach based on the fact that temporal logic verification is computationally cheaper than synthesis. We present a case-study where we compare our approach to the classical non-incremental approach in terms of computation time and memory usage.

I. INTRODUCTION

Temporal logics [1], such as Linear Temporal Logic (LTL) and Computation Tree Logic (CTL), are traditionally used for verification of non-deterministic and probabilistic systems [2]. Even though temporal logics are suitable for specifying complex missions for control systems, they did not gain popularity in the control community until recently [3], [4], [5].

The existing works on control synthesis focus on specifications given in linear time temporal logic. The systems, which sometimes are obtained through an additional abstraction process [3], [6], have finitely many states. With few exceptions [7], their states are fully observable. For such systems, control strategies can be synthesized through exhaustive search of the state space. If the system is deterministic, model checking tools can be easily adapted to generate control strategies [4]. If the system is non-deterministic, the control problem can be mapped to the solution of a Rabin game [8], [6], or a simpler Büchi [9] or GR(1) game [10], if the specification is restricted to fragments of LTL. For probabilistic systems, the LTL control synthesis problem reduces to computing a control policy for a Markov Decision Process (MDP) [11], [12], [13].

In this work, we consider mission specifications expressed as syntactically co-safe LTL formulas [14]. We focus on a particular type of a multi-agent system formed by a deterministically controlled plant and a set of independent, probabilistic, uncontrollable agents, operating on a common, graph-like environment. An illustrative example is a car (plant) approaching a pedestrian crossing, while there are some pedestrians (agents) waiting to cross or already crossing the road. As the state space of the system grows exponentially with the number of pedestrians, one may not be able to utilize any of the existing approaches under computational resource constraints when there is a large number of pedestrians.

We partially address this problem by proposing an incremental control synthesis method that exploits the independence between the components of the system, *i.e.*, the plant modeled as a deterministic transition system and the agents, modeled as Markov chains, and the fact that verification is computationally cheaper than synthesis. We aim to synthesize a plant control strategy that maximize the probability of satisfying a mission specification given as a syntactically co-safe LTL formula. Our method initially considers a considerably smaller agent subset and synthesizes a control policy that maximizes the probability of satisfying the mission specification for the subsystem formed by the plant and this subset. This control policy is then verified against the remaining agents. At each iteration, we remove transitions and states that are not needed in subsequent iterations. This leads to a significant reduction in computation time and memory usage. It is important to note that our method does not need to run to completion. A sub-optimal control policy can be obtained by forcing termination at a given iteration if the computation is performed under stringent resource constraints. It must also be noted that our framework easily extends to the case when the plant is a Markov Decision Process, and we consider a deterministic plant only for simplicity of presentation. We experimentally evaluate the performance of our approach and show that our method clearly outperforms existing non-incremental approaches. Various methods that also use verification during incremental synthesis have been previously proposed in [15], [16]. However, the approach that we present in this paper is, to the best of our knowledge, the first use of verification guided incremental synthesis in the context of probabilistic systems.

The rest of the paper is organized as follows: In Sec. II, we give necessary definitions and some preliminaries in formal methods. The control synthesis problem is formally stated in Sec. III and the solution is presented in Sec. IV. Experimental results are included in Sec. V. We conclude with final remarks in Sec. VI.

II. PRELIMINARIES

For a set Σ, we use $|\Sigma|$ and 2^Σ to denote its cardinality and power set, respectively. A (finite) word ω over a set Σ is

a sequence of symbols $\omega = \omega^0 \ldots \omega^l$ such that $\omega^i \in \Sigma \; \forall i = 0, \ldots, l$.

Definition II.1 (Transition System). *A transition system (TS) is a tuple* $\mathbf{T} := (\mathcal{Q}_\mathrm{T}, q_\mathrm{T}^0, \mathcal{A}_\mathrm{T}, \alpha_\mathrm{T}, \delta_\mathrm{T}, \Pi_\mathrm{T}, \mathcal{L}_\mathrm{T})$, *where*

- $\mathcal{Q}_\mathrm{T}$ *is a finite set of states;*
- $q_\mathrm{T}^0 \in \mathcal{Q}_\mathrm{T}$ *is the initial state;*
- $\mathcal{A}_\mathrm{T}$ *is a finite set of actions;*
- $\alpha_\mathrm{T} : \mathcal{Q}_\mathrm{T} \to 2^{\mathcal{A}_\mathrm{T}}$ *is a map giving the set of actions available at a state;*
- $\delta_\mathrm{T} \subseteq \mathcal{Q}_\mathrm{T} \times \mathcal{A}_\mathrm{T} \times \mathcal{Q}_\mathrm{T}$ *is the transition relation;*
- Π_T *is a finite set of atomic propositions;*
- $\mathcal{L}_\mathrm{T} : \mathcal{Q}_\mathrm{T} \to 2^{\Pi_\mathrm{T}}$ *is a satisfaction map giving the set of atomic propositions satisfied at a state.*

Definition II.2 (Markov Chain). *A (discrete-time, labelled) Markov chain (MC) is a tuple* $\mathbf{M} := (\mathcal{Q}_\mathrm{M}, q_\mathrm{M}^0, \delta_\mathrm{M}, \Pi_\mathrm{M}, \mathcal{L}_\mathrm{M})$, *where* $\mathcal{Q}_\mathrm{M}$, Π_M, *and* $\mathcal{L}_\mathrm{M}$ *are the set of states, the set of atomic propositions, and the satisfaction map, respectively, as in Def. II.1, and*

- $q_\mathrm{M}^0 \in \mathcal{Q}_\mathrm{M}$ *is the initial state;*
- $\delta_\mathrm{M} : \mathcal{Q}_\mathrm{M} \times \mathcal{Q}_\mathrm{M} \to [0, 1]$ *is the transition probability function that satisfies* $\sum_{q' \in \mathcal{Q}_\mathrm{M}} \delta(q, q') = 1 \; \forall q \in \mathcal{Q}_\mathrm{M}$.

In this paper, we are interested in temporal logic missions over a finite time horizon and we use syntactically co-safe LTL formulas [17] to specify them. Informally, a syntactically co-safe LTL formula over the set Π of atomic propositions comprises boolean operators $\neg$ (negation), $\vee$ (disjunction) and $\wedge$ (conjunction), and temporal operators $\mathbf{X}$ (next), $\mathcal{U}$ (until) and $\mathbf{F}$ (eventually). Any syntactically co-safe LTL formula can be written in positive normal form, where the negation operator $\neg$ occurs only in front of atomic propositions. For instance, $\mathbf{X} \, \mathrm{p}$ states that at the next position of the word, proposition p is true. The formula $\mathrm{p}_1 \, \mathcal{U} \, \mathrm{p}_2$ states that there is a future position of the word when proposition p_2 is true, and proposition p_1 is true at least until p_2 is true. For any syntactically co-safe LTL formula ϕ over a set Π, one can construct a FSA with input alphabet 2^Π accepting all and only finite words over 2^Π that satisfy ϕ, which is defined next.

Definition II.3 (Finite State Automaton). *A (deterministic) finite state automaton (FSA) is a tuple* $\mathbf{F} := (\mathcal{Q}_\mathrm{F}, q_\mathrm{F}^0, \Sigma_\mathrm{F}, \delta_\mathrm{F}, \mathcal{F}_\mathrm{F})$, *where*

- $\mathcal{Q}_\mathrm{F}$ *is a finite set of states;*
- $q_\mathrm{F}^0 \in \mathcal{Q}_\mathrm{F}$ *is the initial state;*
- Σ_F *is an input alphabet;*
- $\delta_\mathrm{F} : \mathcal{Q}_\mathrm{F} \times \Sigma_\mathrm{F} \times \mathcal{Q}_\mathrm{F}$ *is a deterministic transition relation;*
- $\mathcal{F}_\mathrm{F} \subseteq \mathcal{Q}_\mathrm{F}$ *is a set of accepting (final) states.*

A *run* of $\mathbf{F}$ over an input word $\omega = \omega^0 \omega^1 \ldots \omega^l$ where $\omega^i \in \Sigma_\mathrm{F} \; \forall i = 0 \ldots l$ is a sequence $r_\mathrm{F} = q^0 q^1 \ldots q^l q^{l+1}$, such that $(q^i, \omega^i, q^{i+1}) \in \delta_\mathrm{F} \; \forall i = 0 \ldots l$ and $q^0 = q_\mathrm{F}^0$. An FSA $\mathbf{F}$ accepts a word over Σ_F if and only the corresponding run ends in some $q \in \mathcal{F}_\mathrm{F}$.

Definition II.4 (Markov Decision Process). *A Markov decision process (MDP) is a tuple* $\mathbf{P} := (\mathcal{Q}_\mathrm{P}, q_\mathrm{P}^0, \mathcal{A}_\mathrm{P}, \alpha_\mathrm{P}, \delta_\mathrm{P}, \Pi_\mathrm{P}, \mathcal{L}_\mathrm{P})$, *where*

- $\mathcal{Q}_\mathrm{P}$ *is a finite set of states;*
- $q_\mathrm{P}^0 \in \mathcal{Q}_\mathrm{P}$ *is the initial state;*
- $\mathcal{A}_\mathrm{P}$ *is a finite set of actions;*
- $\alpha_\mathrm{P} : \mathcal{Q}_\mathrm{P} \to 2^{\mathcal{A}_\mathrm{P}}$ *is a map giving the set of actions available at a state;*
- $\delta_\mathrm{P} : \mathcal{Q}_\mathrm{P} \times \mathcal{A}_\mathrm{P} \times \mathcal{Q}_\mathrm{P} \to [0, 1]$ *is the transition probability function that satisfies* $\sum_{q' \in \mathcal{Q}_\mathrm{P}} \delta(q, a, q') = 1 \; \forall q \in \mathcal{Q}_\mathrm{P}, a \in \alpha_\mathrm{P}(q)$ *and* $\sum_{q' \in \mathcal{Q}_\mathrm{P}} \delta(q, a, q') = 0 \; \forall q \in \mathcal{Q}_\mathrm{P}, a \notin \alpha_\mathrm{P}(q)$.
- Π_P *is a finite set of atomic propositions;*
- $\mathcal{L}_\mathrm{P} : \mathcal{Q}_\mathrm{P} \to 2^{\Pi_\mathrm{P}}$ *is a map giving the set of atomic propositions satisfied in a state.*

For an MDP $\mathbf{P}$, we define a stationary policy $\mu_\mathrm{P} : \mathcal{Q}_\mathrm{P} \to \mathcal{A}_\mathrm{P}$ such that for a state $q \in \mathcal{Q}_\mathrm{P}$, $\mu_\mathrm{P}(q) \in \alpha_\mathrm{P}(q)$. This stationary policy can then be used to resolve all nondeterministic choices in $\mathbf{P}$ by applying action $\mu(q)$ at each $q \in \mathcal{Q}_\mathrm{P}$. A path of $\mathbf{P}$ under policy μ_P is a finite sequence of states $r_\mathrm{P}^{\mu_\mathrm{P}} = q^0 q^1 \ldots q^l$ such that $l \geq 0$, $q^0 = q_\mathrm{P}^0$ and $\delta_\mathrm{P}(q^{k-1}, \mu_\mathrm{P}(q^{k-1}), q^k) > 0 \; \forall k \in [1, l]$. A path $r_\mathrm{P}^{\mu_\mathrm{P}}$ generates a finite word $\mathcal{L}_\mathrm{P}(r_\mathrm{P}^{\mu_\mathrm{P}}) = \mathcal{L}_\mathrm{P}(q^0) \mathcal{L}_\mathrm{P}(q^1) \ldots \mathcal{L}_\mathrm{P}(q^l)$ where $\mathcal{L}_\mathrm{P}(q^k)$ is the set of atomic propositions satisfied at state q^k. Next, we use $Paths_\mathrm{P}^{\mu_\mathrm{P}}$ to denote the set of all paths of $\mathbf{P}$ under a policy μ_P. Finally, we define $P^{\mu_\mathrm{P}}(\phi)$ as the probability of satisfying ϕ under policy μ_P.

Remark II.5. *Syntactically co-safe LTL formulas have infinite time semantics, thus they are actually interpreted over infinite words [17]. Measurability of languages satisfying LTL formulas is also defined for infinite words generated by infinite paths [2]. However, one can determine whether a given infinite word satisfies a syntactically co-safe LTL formula by considering only a finite prefix of it. It can be easily shown that our above definition of* $Paths_\mathrm{P}^{\mu_\mathrm{P}}$ *inherits the same measurability property given in [2].*

III. PROBLEM FORMULATION AND APPROACH

In this section we introduce the control synthesis problem with temporal constraints for a system that models a plant operating in the presence of probabilistic independent agents.

A. System Model

Consider a system consisting of a deterministic plant that we can control (*e.g.*, a robot) and n agents operating in an environment modeled by a graph $\mathcal{E} = (V, \to_\mathcal{E}, \mathcal{L}_\mathcal{E}, \Pi_\mathcal{E})$, where V is the set of vertices, $\to_\mathcal{E} \subseteq V \times V$ is the set of edges, and $\mathcal{L}_\mathcal{E}$ is the labeling function that maps each vertex to a proposition in $\Pi_\mathcal{E}$. For example, $\mathcal{E}$ can be the quotient graph of a partitioned environment, where V is a set of labels for the regions in the partition and $\to_\mathcal{E}$ is the corresponding adjacency relation (see Figs. 1, 2). Agent i is modeled as an MC $\mathbf{M}_i = (\mathcal{Q}_i, q_i^0, \delta_i, \Pi_i, \mathcal{L}_i)$, with $\mathcal{Q}_i \subseteq V$ and $\delta_i \subseteq \to_\mathcal{E}$, $i = 1, \ldots, n$. The plant is assumed to be a deterministic transition system TS $\mathbf{T} = (\mathcal{Q}_\mathrm{T}, q_\mathrm{T}^0, \mathcal{A}_\mathrm{T}, \alpha_\mathrm{T}, \delta_\mathrm{T}, \Pi_\mathrm{T}, \mathcal{L}_\mathrm{T})$, where $\mathcal{Q}_\mathrm{T} \subseteq V$ and $\delta_t \subseteq \to_\mathcal{E}$. We assume that all components of the system (the plant and the agents) make transitions synchronously by picking edges of the graph. We also assume that the state of the system is perfectly known at

Fig. 1: A partitioned road environment, where a car (plant) is required to reach c_4 without colliding with any of the pedestrians (agents).

Fig. 2: TS $\mathbf{T}$ and MCs $\mathbf{M}_1 \ldots \mathbf{M}_5$ that model the car and the pedestrians.

Fig. 3: Deterministic FSA $\mathbf{F}$ that corresponds to $\phi = \neg\texttt{col}\,\mathcal{U}\,\texttt{end}$ where $\texttt{col} = \bigvee_{i=1\ldots5,j=0\ldots4}(c_j^{\mathrm{T}} \wedge c_j^i)$ and $\texttt{end} = c_4^{\mathrm{T}}$. q_0 and q_1 are initial and final states, respectively.

any given instant and we can control the plant but we have no control over the agents.

We define the sets of propositions and labeling functions of the individual components of the system such that they inherit the propositions of their current vertex from the graph while preserving their own identities. Formally, we have $\Pi_{\mathrm{T}} = \{(\mathrm{T}, \mathcal{L}_{\mathcal{E}}(q)) | q \in \mathcal{Q}_{\mathrm{T}}\}$ and $\mathcal{L}_{\mathrm{T}}(q) = (\mathrm{T}, \mathcal{L}_{\mathcal{E}}(q))$ for the plant, and $\Pi_i = \{(i, \mathcal{L}_{\mathcal{E}}(q)) | q \in \mathcal{Q}_i\}$ and $\mathcal{L}_i(q) = (i, \mathcal{L}_{\mathcal{E}}(q))$ for agent i. Finally, we define the set Π of propositions as $\Pi = \Pi_{\mathrm{T}} \cup \Pi_i \cup \ldots \cup \Pi_n \subseteq \{(i,p) | i = \{\mathrm{T}, 0, \ldots, n\}, p \in \Pi_{\mathcal{E}}\}$.

B. Problem Formulation

As it will become clear in Sec. IV-D, the joint behavior of the plant and agents in the graph environment can be modeled by the parallel composition of the TS and MC models described above, which takes the form of an MDP (see Def. II.4). Given a syntactically co-safe LTL formula ϕ over Π, our goal is to synthesize a policy for this MDP, which we will simply refer to as the *system*, such that the probability of satisfying ϕ is either maximized or above a given threshold. Since we assume perfect state information, the plant can implement a control policy computed for the system, i.e, based on its state and the state of all the other agents. As a result, we will not distinguish between a control policy for the plant and a control policy for the system, and we will refer to it simply as *control policy*. We can now formulate the main problem considered in this paper:

Problem III.1. *Given a system described by a plant $\mathbf{T}$ and a set of agents $\mathbf{M}_1, \ldots, \mathbf{M}_n$ operating on a graph $\mathcal{E}$, and given a specification in the form of a syntactically co-safe LTL formula ϕ over Π, synthesize a control policy $\mu^\star$ that satisfies the following objective: (a) If a probability threshold p_{thr} is given, the probability that the system satisfies ϕ under $\mu^\star$ exceeds p_{thr}. (b) Otherwise, $\mu^\star$ maximizes the probability that the system satisfies ϕ. If no such policy exists, report failure.*

As will be shown in Sec. IV-A, the parallel composition of MDP and MC models also takes the form of an MDP. Hence, our approach can easily accommodate the case where

the plant is a Markov Decision Process. We consider a deterministic plant only for simplicity of presentation.

Example III.2. *Fig. 1 illustrates a car in a 5-cell environment with 5 pedestrians, where $\mathcal{L}_{\mathcal{E}}(v) = v$ for $v \in \{c_0, \ldots, c_4\}$. Fig. 2 illustrates the TS $\mathbf{T}$ and the MCs $\mathbf{M}_1, \ldots, \mathbf{M}_5$ that model the car and the pedestrians. The car is required to reach the end of the crossing (c_4) without colliding with any of the pedestrians. To enforce this behavior, we write our specification as*

$$\phi := \left(\neg \bigvee_{i=1\ldots5,j=0\ldots4} ((\mathrm{T}, c_j) \wedge (i, c_j)) \right) \mathcal{U}\,(\mathrm{T}, c_4). \quad (1)$$

The deterministic FSA that corresponds to ϕ is given in Fig. 3, where $\texttt{col} = \bigvee_{i=1\ldots5,j=0\ldots4}((\mathrm{T}, c_j) \wedge (i, c_j))$ and $\texttt{end} = (\mathrm{T}, c_4)$.

C. Solution Outline

One can directly solve Prob. III.1 by reducing it to a Maximal Reachability Probability (MRP) problem on the MDP modeling the overall system [18]. This approach, however, is very resource demanding as it scales exponentially with the number agents. As a result, the environment size and the number of agents that can be handled in a reasonable time frame and with limited memory are small. To address this issue, we propose a highly efficient incremental control synthesis method that exploits the independence between the system components and the fact that verification is less demanding than synthesis. At each iteration i, our method will involve the following steps: synthesis of an optimal control policy considering only some of the agents (Sec. IV-D), verification of this control policy with respect to the complete system (Sec. IV-E) and minimization of the system model under the guidance of this policy (Sec. IV-F).

IV. PROBLEM SOLUTION

Our solution to Prob. III.1 is given in the form of Alg. 1. In the rest of this section, we explain each of its steps in detail.

Algorithm 1: INCREMENTAL-CONTROL-SYNTHESIS

Input: $\mathbf{T}, \mathbf{M}_1, \ldots, \mathbf{M}_n, \phi, (p_{thr})$.
Output: $\mu^\star$ s.t. $P^{\mu^\star}(\phi) \geq P^\mu(\phi) \, \forall \mu$ if p_{thr} is not given, otherwise $P^{\mu^\star}(\phi) > p_{thr}$.

1 $\mathcal{M} \leftarrow \{\mathbf{M}_1, \ldots, \mathbf{M}_n\}$.
2 Construct FSA $\mathbf{F}$ corresponding to ϕ.
3 $\mu^\star \leftarrow \emptyset$, $P_{\mathcal{M}}^{\mu^\star}(\phi) \leftarrow 0$, $\mathcal{M}_0 \leftarrow \emptyset$, $\mathbf{A}_0 \leftarrow \mathbf{T}$, $i \leftarrow 1$.
4 Process ϕ to form $\mathcal{M}_i^{new}$.
5 **while** *True* **do**
6     $\mathcal{M}_i \leftarrow \mathcal{M}_{i-1} \cup \mathcal{M}_i^{new}$.
7     $\mathbf{A}_i \leftarrow \mathbf{A}_{i-1} \otimes \mathcal{M}_i^{new}$.
8     $\mathbf{P}_i \leftarrow \mathbf{A}_i \otimes \mathbf{F}$.
9   Synthesize μ_i that maximizes $P_{\mathcal{M}_i}^{\mu_i}(\phi)$ using $\mathbf{P}_i$.
10   **if** p_{thr} *given* **then**
11   **if** $P_{\mathcal{M}_i}^{\mu_i}(\phi) < p_{thr}$ **then**
12   Fail: $\nexists \mu$ such that $P^\mu(\phi) \geq p_{thr}$.
13   **else if** $\mathcal{M}_i = \mathcal{M}$ **then**
14   Success: $\mu^\star \leftarrow \mu_i$, Return $\mu^\star$.
15   **else**
16   Continue with verification on line 20.
17   **else if** $\mathcal{M}_i = \mathcal{M}$ **then**
18   Success: $\mu^\star \leftarrow \mu_i$, Return $\mu^\star$.
19   **else**
20   Obtain the MC $\mathbf{M}_{\mathcal{M}_i}^{\mu_i}$ induced on $\mathbf{P}_i$ by μ_i.
21       $\overline{\mathcal{M}_i} \leftarrow \mathcal{M} \setminus \mathcal{M}_i$.
22       $\mathbf{M}_{\mathcal{M}}^{\mu_i} \leftarrow \mathbf{M}_{\mathcal{M}_i}^{\mu_i} \otimes \overline{\mathcal{M}_i}$.
23       $\mathbf{V}_i \leftarrow \mathbf{M}_{\mathcal{M}}^{\mu_i} \otimes \mathbf{F}$.
24   Compute $P_{\mathcal{M}}^{\mu_i}(\phi)$ using $\mathbf{V}_i$.
25   **if** $P_{\mathcal{M}}^{\mu_i}(\phi) > P_{\mathcal{M}}^{\mu^\star}(\phi)$ **then**
26          $\mu^\star \leftarrow \mu_i$, $P_{\mathcal{M}}^{\mu^\star}(\phi) \leftarrow P_{\mathcal{M}}^{\mu_i}(\phi)$.
27   **if** p_{thr} *given and* $P_{\mathcal{M}}^{\mu^\star}(\phi) > p_{thr}$ **then**
28   Success: Return $\mu^\star$.
29   **else**
30   Set $\mathcal{M}_{i+1}^{new}$ to some agent $\mathbf{M}_j \in \overline{\mathcal{M}_i}$.
31   Minimize $\mathbf{A}_i$.
32   Increment i.

A. Parallel Composition of System Components

Given the set $\mathcal{M} = \{\mathbf{M}_1, \ldots, \mathbf{M}_n\}$ of all agents, we use $\mathcal{M}_i \subseteq \mathcal{M}$ to denote its subset used at iteration i. Then, we define the synchronous parallel composition $\mathbf{T} \otimes \mathcal{M}_i$ of $\mathbf{T}$ and agents in $\mathcal{M}_i = \{\mathbf{M}_{i1}, \ldots, \mathbf{M}_{ij}\}$ for different types of $\mathbf{T}$ as follows.

If $\mathbf{T}$ is a TS, then we define $\mathbf{T} \otimes \mathcal{M}_i$ as the MDP $\mathbf{A} = (\mathcal{Q}_A, q_A^0, \mathcal{A}_A, \alpha_A, \delta_A, \Pi_A, \mathcal{L}_A) = \mathbf{T} \otimes \mathcal{M}_i$, such that

- $\mathcal{Q}_A \subseteq \mathcal{Q}_T \times \mathcal{Q}_{i1} \times \ldots \times \mathcal{Q}_{ij}$ such that a state $q = (q_T, q_{i1}, \ldots, q_{ij})$ exists iff it is reachable from the initial states;
- $q_A^0 = (q_T^0, q_{i1}^0, \ldots, q_{ij}^0)$;
- $\mathcal{A}_A = \mathcal{A}_T$;
- $\alpha_A(q) = \alpha_T(q_T)$, where q_T is the element of q that corresponds to the state of $\mathbf{T}$;
- $\Pi_A = \Pi_T \cup \Pi_{i1} \cup \ldots \cup \Pi_{ij}$;
- $\mathcal{L}_A(q) = \mathcal{L}_T(q_T) \cup \mathcal{L}_{i1}(q_{i1}) \cup \ldots \cup \mathcal{L}_{ij}(q_{ij})$;
- $\delta_A(q = (q_T, q_{i1}, \ldots, q_{ij}), a, q' = (q_T', q_{i1}', \ldots, q_{ij}')) = 1\{(q_T, a, q_T') \in \delta_T\} \times \delta(q_{i1}, q_{i1}') \times \ldots \times \delta(q_{ij}, q_{ij}')$,

where $1\{\cdot\}$ is the indicator function.

If $\mathbf{T}$ is an MDP, then we define $\mathbf{T} \otimes \mathcal{M}_i$ as the MDP $\mathbf{A} = (\mathcal{Q}_A, q_A^0, \mathcal{A}_A, \alpha_A, \delta_A, \Pi_A, \mathcal{L}_A) = \mathbf{T} \otimes \mathcal{M}_i$, such that $\mathcal{Q}_A, q_A^0, \mathcal{A}_A, \alpha_A, \Pi_A$, and $\mathcal{L}_A$ are as given in the case where $\mathbf{T}$ is a TS and

- $\delta_A(q = (q_T, q_{i1}, \ldots, q_{ij}), a, q' = (q_T', q_{i1}', \ldots, q_{ij}')) = \delta_T(q_T, a, q_T') \times \delta_{i1}(q_{i1}, q_{i1}') \times \ldots \times \delta_{ij}(q_{ij}, q_{ij}')$.

Finally if $\mathbf{T}$ is an MC, then we define $\mathbf{T} \otimes \mathcal{M}_i$ as the MC $\mathbf{A} = (\mathcal{Q}_A, q_A^0, \delta_A, \Pi_A, \mathcal{L}_A) = \mathbf{T} \otimes \mathcal{M}_i$ where $\mathcal{Q}_A, q_A^0, \Pi_A, \mathcal{L}_A$ are as given in the case where $\mathbf{T}$ is a TS and

- $\delta_A(q = (q_T, q_{i1}, \ldots, q_{ij}), q' = (q_T', q_{i1}', \ldots, q_{ij}')) = \delta_T(q_T, q_T') \times \delta_{i1}(q_{i1}, q_{i1}') \times \ldots \times \delta_{ij}(q_{ij}, q_{ij}')$.

B. Product MDP and Product MC

Given the deterministic FSA $\mathbf{F}$ that recognizes all and only the finite words that satisfy ϕ, we define the product of $\mathbf{M} \otimes \mathbf{F}$ for different types of $\mathbf{M}$ as follows.

If $\mathbf{M}$ is an MDP, we define $\mathbf{M} \otimes \mathbf{F}$ as the product MDP $\mathbf{P} = (\mathcal{Q}_P, q_P^0, \mathcal{A}_P, \alpha_P, \delta_P, \Pi_P, \mathcal{L}_P) = \mathbf{M} \otimes \mathbf{F}$, where

- $\mathcal{Q}_P \subseteq \mathcal{Q}_M \times \mathcal{Q}_F$ such that a state q exists iff it is reachable from the initial states;
- $q_P^0 = (q_M^0, q_F)$ such that $(q_F^0, \mathcal{L}_M(q_M), q_F) \in \delta_F$;
- $\mathcal{A}_P = \mathcal{A}_M$;
- $\alpha_P((q_M, q_F)) = \alpha_M(q_M)$;
- $\Pi_P = \Pi_M$;
- $\mathcal{L}_P((q_M, q_F)) = \mathcal{L}_M(q_M)$;
- $\delta_P((q_M, q_F), a, (q_M', q_F')) = 1\{(q_F, \mathcal{L}_M(q_M'), q_F') \in \delta_F\} \times \delta_M(q_M, a, q_M')$,

where $1\{\cdot\}$ is the indicator function. In this product MDP, we also define the set $\mathcal{F}_P$ of final states such that a state $q = (q_M, q_F) \in \mathcal{F}_P$ iff $q_F \in \mathcal{F}_F$, where $\mathcal{F}_F$ is the set of final states of $\mathbf{F}$.

If $\mathbf{M}$ is an MC, we define $\mathbf{M} \otimes \mathbf{F}$ as the product MC $\mathbf{P} = (\mathcal{Q}_P, q_P^0, \delta_P, \Pi_P, \mathcal{L}_P) = \mathbf{M} \otimes \mathbf{F}$ where $\mathcal{Q}_P, q_P^0, \Pi_P, \mathcal{L}_P$ are as given in the case where $\mathbf{M}$ is an MDP and

- $\delta_P((q_M, q_F), (q_M', q_F')) = 1\{(q_F, \mathcal{L}_M(q_M'), q_F') \in \delta_F\} \times \delta_M(q_M, q_M')$.

In this product MC, we also define the set $\mathcal{F}_P$ of final states as given above.

C. Initialization

Lines 1 to 4 of Alg. 1 correspond to the initialization procedure of our algorithm. First, we form the set $\mathcal{M} = \{\mathbf{M}_1, \ldots, \mathbf{M}_n\}$ of all agents and construct the FSA $\mathbf{F}$ that corresponds to ϕ. Such $\mathbf{F}$ can be automatically constructed

using existing tools, *e.g.*, [19]. Since we have not synthesized any control policies so far, we reset the variable $\mu^\star$ that holds the best policy at any given iteration and set the probability $P_\mathcal{M}^{\mu^\star}(\phi)$ of satisfying ϕ under policy $\mu^\star$ in the presence of agents in $\mathcal{M}$ to 0. As we have not considered any agents so far, we set the subset $\mathcal{M}_0$ to be an empty set. We then set $\mathbf{A}_0$, which stands for the parallel composition of the plant $\mathbf{T}$ and the agents in $\mathcal{M}_0$, to $\mathbf{T}$. We also initialize the iteration counter i to 1.

Line 4 of Alg. 1 initializes the set $\mathcal{M}_1^{new}$ of agents that will be considered in the synthesis step of the first iteration of our algorithm. In order to be able to guarantee completeness, we require this set to be the maximal set of agents that satisfy the mission, *i.e.*, the agent subset that can satisfy ϕ but not strictly needed to satisfy ϕ. To form $\mathcal{M}_1^{new}$, we first rewrite ϕ in positive normal form to obtain ϕ_{pnf}, where the negation operator $\neg$ occurs only in front of atomic propositions. Conversion of ϕ to ϕ_{pnf} can be performed automatically using De Morgan's laws and equivalences for temporal operators as given in [2]. Then, using this fact, we include an agent $\mathbf{M}_i \in \mathcal{M}$ in $\mathcal{M}_1^{new}$ if any of its corresponding propositions of the form $(i, p), p \in \Pi_i$ appears non-negated in ϕ_{pnf}. For instance, given $\phi := \neg((3, \texttt{p3}) \wedge (\mathbf{T}, \texttt{p3}))\, \mathcal{U}\, ((1, \texttt{p1}) \vee (2, \texttt{p2}))$, either one of agents $\mathbf{M}_1$ and $\mathbf{M}_2$ can satisfy the formula, whereas agent $\mathbf{M}_3$ can only violate it. Therefore, for this example we set $\mathcal{M}_1^{new} = \{\mathbf{M}_1, \mathbf{M}_2\}$. In case $\mathcal{M}_1^{new} = \emptyset$ after this procedure, we form $\mathcal{M}_1^{new}$ arbitrarily by including some agents from $\mathcal{M}$ and proceed with the synthesis step of our approach.

D. Synthesis

Lines 6 to 19 of Alg. 1 correspond to the synthesis step of our algorithm. At the i^{th} iteration, the agent subset that we consider is given by $\mathcal{M}_i = \mathcal{M}_{i-1} \cup \mathcal{M}_i^{new}$ where $\mathcal{M}_i^{new}$ contains the agents that will be newly considered as provided by the previous iteration's verification stage or by the initialization procedure given in Sec. IV-C if i is 1. First, we construct the parallel composition $\mathbf{A}_i = \mathbf{A}_{i-1} \otimes \mathcal{M}_i^{new}$ of our plant and the agents in $\mathcal{M}_i$ as described in Sec. IV-A. Notice that, we use $\mathbf{A}_{i-1}$ to save from computation time and memory as $\mathbf{A}_{i-1} \otimes \mathcal{M}_i^{new}$ is typically smaller than $\mathbf{T} \otimes \mathcal{M}_i$ due to the minimization procedure explained in Sec. IV-F. Next, we construct the product MDP $\mathbf{P}_i = \mathbf{A}_i \otimes \mathbf{F}$ as explained in Sec. IV-B. Then, our control synthesis problem can be solved by solving a maximal reachability probability (MRP) problem on $\mathbf{P}_i$ where one computes the maximum probability of reaching the set $\mathcal{F}_\mathbf{P}$ from the initial state $q_\mathbf{P}^0$ [18], after which the corresponding optimal control policy μ_i can be recovered as given in [2], [13]. Consequently, at line 9 of Alg. 1 we solve the MRP problem on $\mathbf{P}_i$ using value iteration to obtain optimal policy μ_i that maximizes the probability of satisfaction of ϕ in the presence of the agents in $\mathcal{M}_i$. We denote this probability by $P_{\mathcal{M}_i}^{\mu_i}(\phi)$, whereas $P^\mu(\phi)$ stands for the probability that the complete system satisfies ϕ under policy μ.

The steps that we take at the end of the synthesis, *i.e.*, lines 10 to 19 of Alg. 1, depends on whether p_{thr} is given or not.

At any iteration i, if p_{thr} is given and $P_{\mathcal{M}_i}^{\mu_i}(\phi) < p_{thr}$, we terminate by reporting that there exists no control policy $\mu : P^\mu(\phi) \geq p_{thr}$ which is a direct consequence of Prop. IV.1. If p_{thr} is given and $P_{\mathcal{M}_i}^{\mu_i}(\phi) \geq p_{thr}$, we consider the following cases. If $\mathcal{M}_i = \mathcal{M}$, we set $\mu^\star$ to μ_i and return $\mu^\star$ as it satisfies the probability threshold. Otherwise, we proceed with the verification of μ_i as there are remaining agents that were not considered during synthesis and can potentially violate ϕ. For the case where p_{thr} is not given we consider the current agent subset $\mathcal{M}_i$. If $\mathcal{M}_i = \mathcal{M}$ we terminate and return $\mu^\star$ as there are no agents left to consider. Otherwise, we proceed with the verification stage.

Proposition IV.1. *The sequence $\{P_{\mathcal{M}_i}^{\mu_i}(\phi)\}$ is non-increasing.*

Proof. As given in Sec. IV-C, $\mathcal{M}_1$ includes all those agents that can satisfy the propositions that lead to satisfaction of ϕ. Let $pref(\phi)$ be the set of finite words that satisfy ϕ and let MC $\mathbf{M}_j$ of agent j be such that $\mathbf{M}_j \notin \mathcal{M}_1$. Consider a finite satisfying word σ such that $\sigma = \sigma^0 \sigma^1 \ldots \sigma^l \in pref(\phi)$. Suppose there exists an index $k \in \{0, \ldots, l\}$ such that for some $q \in \mathcal{Q}_j$ and $\mathcal{L}_j(q) \in \sigma^k$. Then, $\tilde{\sigma} = \sigma^0 \sigma^1 \ldots \sigma^{k-1} \tilde{\sigma}^k \sigma^{k+1} \ldots \sigma^l$ is also in $pref(\phi)$ where $\tilde{\sigma}^k = \sigma^k \setminus \mathcal{L}_j(q)$. Now, let $r = q^0 q^1 \ldots q^l$ be a path of the system after including $\mathbf{M}_j$. Let $\omega = \mathcal{L}(r) = \omega^0 \omega^1 \ldots \omega^l$ be the word generated by r. If ω satisfies ϕ, then $\tilde{\omega} = \tilde{\omega}^0 \tilde{\omega}^1 \ldots \tilde{\omega}^l$ also satisfies ϕ where $\tilde{\omega}^k = \omega^k \setminus \mathcal{L}_j(q_j^k)$ for each $k \in \{0, \ldots, l\}$ and q_j^k is the state of $\mathbf{M}_j$ in q^k. Thus, we conclude that the set of paths that satisfy ϕ cannot increase after we add agent $\mathbf{M}_j \in \mathcal{M} \setminus \mathcal{M}_1$, and the sequence $\{P_{\mathcal{M}_i}^{\mu_i}(\phi)\}$ is non-increasing such that it attains its maximum value $P_{\mathcal{M}_1}^{\mu_1}(\phi)$ at the first iteration and does not increase as more agents from $\mathcal{M} \setminus \mathcal{M}_1$ are considered in the following iterations. ∎

Corollary IV.2. *If at any iteration $P_{\mathcal{M}_i}^{\mu_i}(\phi) < p_{thr}$, then there does not exist a policy $\mu : P^\mu(\phi) \geq p_{thr}$, where μ_i is an optimal control policy that we compute at the synthesis stage of the i^{th} iteration considering only the agents in $\mathcal{M}_i$.*

E. Verification and Selection of $\mathcal{M}_{i+1}^{new}$

Lines 20 to 30 of Alg. 1 correspond to the verification stage of our algorithm. In the verification stage, we verify the policy μ_i that we have just synthesized considering the entire system and accordingly update the best policy so far, which we denote by $\mu^\star$.

Note that μ_i maximizes the probability of satisfying ϕ in the presence of agents in $\mathcal{M}_i$ and induces an MC by resolving all non-deterministic choices in $\mathbf{P}_i$. Thus, we first obtain the induced Markov Chain $\mathbf{M}_{\mathcal{M}_i}^{\mu_i}$ that captures the joint behavior of the plant and the agents in $\mathcal{M}_i$ under policy μ_i. Then, we proceed by considering the agents that were not considered during synthesis of μ_i, *i.e.*, agents in $\overline{\mathcal{M}_i} = \mathcal{M} \setminus \mathcal{M}_i$. In order to account for the existence of the agents that we newly consider, we exploit the independence between the systems and construct the MC $\mathbf{M}_\mathcal{M}^{\mu_i} = \mathbf{M}_{\mathcal{M}_i}^{\mu_i} \otimes \overline{\mathcal{M}_i}$ in line 22. In lines 23 and 24 of Alg. 1, we construct the product MC $\mathbf{V}_i = \mathbf{M}_\mathcal{M}^{\mu_i} \otimes \mathbf{F}$ and compute the probability $P_\mathcal{M}^{\mu_i}(\phi)$ of

satisfying ϕ in the presence of all agents in $\mathcal{M}$ by computing the probability of reaching $\mathbf{V}_i$'s final states from its initial state using value iteration. Finally, in lines 25 and 26 we update $\mu^\star$ so that $\mu^\star = \mu_i$ if $P_\mathcal{M}^{\mu_i}(\phi) > P_\mathcal{M}^{\mu^\star}(\phi)$, i.e., if we have a policy that is better than the best we have found so far. Notice that, keeping track of the best policy $\mu^\star$ makes Alg. 1 an anytime algorithm, i.e., the algorithm can be terminated as soon as some $\mu^\star$ is obtained.

At the end of the verification stage, if p_{thr} is given and $P_\mathcal{M}^{\mu^\star}(\phi) \geq p_{thr}$ we terminate and return $\mu^\star$, as it satisfies the given probability threshold. Otherwise in line 30 of Alg. 1, we pick an arbitrary $\mathbf{M}_j \in \overline{\mathcal{M}}_i$ to be included in $\mathcal{M}^{i+1}$, which we call the *random agent first (RAF)* rule. Note that, one can also choose to pick the smallest $\mathbf{M}_j$ in terms of state and transition count to minimize the overall computation time, which we call the *smallest agent first (SAF)* rule.

Proposition IV.3. *The sequence $\{P_\mathcal{M}^{\mu^\star}(\phi)\}$ is a non-decreasing sequence.*

Proof. The result directly follows from the fact that $\mu^\star$ is set to μ_i if and only if $P_\mathcal{M}^{\mu_i}(\phi) > P_\mathcal{M}^{\mu^\star}(\phi)$. ∎

F. Minimization

The minimization stage of our approach (line 31 in Alg. 1) aims to reduce the overall resource usage by removing those transitions and states of $\mathbf{A}_i$ that are not needed in the subsequent iterations. We first set the minimization threshold p_{min} to p_{thr} if given, otherwise we set it to $P_\mathcal{M}^{\mu^\star}(\phi)$. Next, we iterate over the states of $\mathbf{P}_i$ and check the maximum probability of satisfying the mission under each available action. Note that, the value iteration that we perform in the synthesis step already provides us with the maximum probability of satisfying ϕ from any state in $\mathbf{P}_i$. Then, we remove an action a from state q_A in $\mathbf{A}_i$ if for all $q_\mathrm{F} \in Q_\mathrm{F}$, the maximum probability of satisfying the mission by taking action a at $(q_\mathrm{A}, q_\mathrm{F})$ in $\mathbf{P}_i$ is below p_{min}. After removing the transitions corresponding to all such actions, we also prune any orphan states in $\mathbf{A}_i$, i.e., states that are not reachable from the initial state. Then, we proceed with the synthesis stage of the next iteration.

Proposition IV.4. *Minimization phase does not affect the correctness and the completeness of our approach.*

Proof. To prove the correctness, we need to show that for an arbitrary policy μ on the minimized MDP $\mathbf{A}_{min}$, the probability that $\mathbf{A}_{min}$ satisfies ϕ under μ is equal to the probability that $\mathbf{A}$ satisfies ϕ under μ where $\mathbf{A}$ is the original MDP before minimization. Correctness, in this case, follows directly from the fact that, in each state q, we do not modify the transition probabilities associated with an action that is enabled in q after minimization. Thus, it remains to show that minimization does not affect the completeness of the approach. We first consider the removal of orphaned states. Since these states cannot be reached from the initial state, they also will not be a part of any feasible control policy, and their removal does not affect the completeness of the approach. Finally, we consider the removal of those

actions that drive the system to the set of target states with probability smaller than the minimization threshold. For the case where we use p_{thr}, completeness is not affected as we remove only those transitions that we would not take as we are looking for control policies with $P_\mathcal{M}^\mu(\phi) \geq p_{thr}$ and $P_{\mathcal{M}_i}^{\mu_i}$ is a non-increasing sequence (Prop. IV.1). For the case where we use $P_\mathcal{M}^{\mu^\star}$, we also remove those transitions that we would not take as $P_\mathcal{M}^{\mu^\star}$ is a non-decreasing sequence (Prop. IV.3). Hence, the minimization procedure does not affect the completeness of the overall approach as well. ∎

We finally show that Alg. 1 correctly solves Prob. III.1.

Proposition IV.5. *Alg. 1 solves Prob. III.1.*

Proof. Alg. 1 combines all the steps given in this section and synthesizes a control policy $\mu^\star$ that either ensures $P^{\mu^\star}(\phi) \geq p_{thr}$ if p_{thr} is given, or maximizes $P^{\mu^\star}(\phi)$. If Alg. 1 terminates in line 12, completeness is guaranteed by the fact that $P_{\mathcal{M}_i}^{\mu_i}$ is a non-increasing sequence as given in Prop. IV.1. Also, as given in Prop. IV.4, minimization stage does not affect the correctness and completeness of the approach. Thus, Alg. 1 solves Prob. III.1. ∎

V. EXPERIMENTAL RESULTS

In this section we return to the pedestrian crossing problem given in Example III.2 and illustrated in Figs. 1, 2. The mission specification ϕ for this example is given in Eq. (1). In the following, we compare the performance of our incremental algorithm with the performance of the classical method that attempts to solve this problem in a single pass using value iteration as in [18].

In our experiments we used an iMac i5 quad-core desktop computer and considered C++ implementations of both approaches. During the experiments, our algorithm picked the new agent $\mathcal{M}_i^{new}$ to be considered at the next iteration in the following order: $\mathbf{M}_1, \mathbf{M}_2, \mathbf{M}_3, \mathbf{M}_4, \mathbf{M}_5$, i.e., according to the smallest agent first rule given in Sec. IV-E.

When no p_{thr} was given, optimal control policies synthesized by both of the algorithms satisfied ϕ with a probability of 0.8. The classical approach solved the control synthesis problem in 6.75 seconds, and the product MDP on which the MRP problem was solved had 1004 states and 26898 transitions. In comparison, our incremental approach solved the same problem in 4.44 seconds, thanks to the minimization stage of our approach, which reduced the size of the problem at every iteration by pruning unneeded actions and states. The largest product MDP on which the MRP problem was solved in the synthesis stage of our approach had 266 states and 4474 transitions. The largest product MC that was considered in the verification stage of our approach had 405 states and 6125 transitions. The probabilities of satisfying ϕ under policy μ_i obtained at each iteration of our algorithm were $P_\mathcal{M}^{\mu_1}(\phi) = 0.463$, $P_\mathcal{M}^{\mu_2}(\phi) = 0.566$, $P_\mathcal{M}^{\mu_3}(\phi) = 0.627$, $P_\mathcal{M}^{\mu_4}(\phi) = 0.667$, and $P_\mathcal{M}^{\mu_5}(\phi) = 0.8$. When p_{thr} was given as 0.65, our approach finished in 3.63 seconds and terminated after the fourth iteration returning a sub-optimal control policy with a 0.667 probability of satisfying ϕ. In
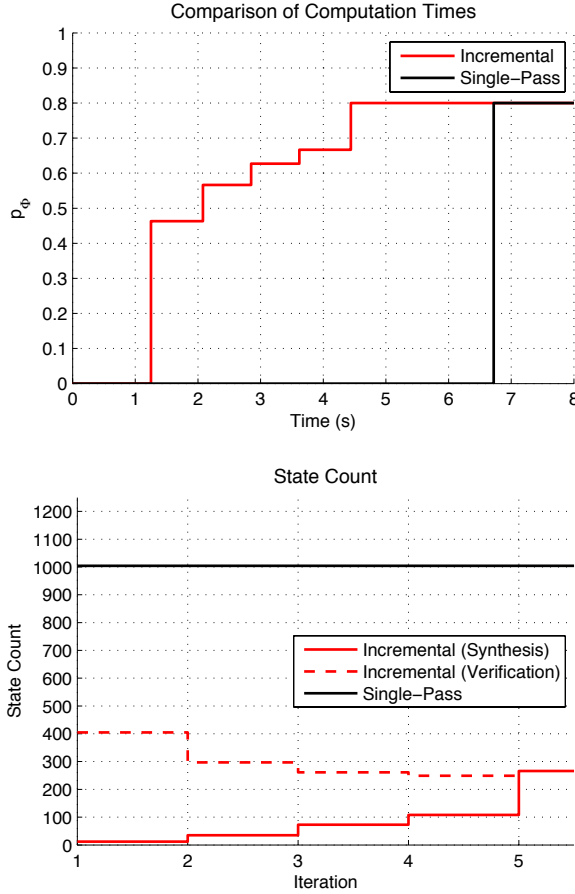
Fig. 4: Comparison of the classical single-pass and proposed incremental algorithms. The top plot shows the running times of the algorithms and the probabilities of satisfying ϕ under synthesized policies. The bottom plot compares the state counts of the product MDPs on which the MRP problem was solved in both approaches (black and red lines) and shows the state count of the product MC considered in the verification stage of our incremental algorithm (red dashed line).

this case, the largest product MDP on which the MRP problem was solved had only 99 states and 680 transitions. Furthermore, since our algorithm runs in an anytime manner, it could be terminated as soon as a control policy was available, *i.e.*, at the end of the first iteration (1.25 seconds). Fig. 4 compares the classical single-pass approach with our incremental algorithm in terms of running time and state counts of the product MDPs and MCs.

It is interesting to note that state count of the product MDP considered in the synthesis stage of our algorithm increases as more agents are considered, whereas state count of the product MC considered in the verification stage of our algorithm decreases as the minimization stage removes unneeded states and transitions after each iteration. It must also be noted that, $|\mathcal{M}_1|$, *i.e.*, cardinality of the initial agent subset, is an important factor for the performance of our algorithm. As discussed in this section, for $|\mathcal{M}_1| << |\mathcal{M}|$ our algorithm outperforms the classical method both in terms of running time and memory usage. However, for $|\mathcal{M}_1| \sim |\mathcal{M}|$ we expect the resource usage of our algorithm

to be close to that of the classical approach, as in this case almost all of the agents will be considered in the synthesis stage of the first iteration. We plan to address this issue in future work. Nevertheless, most typical finite horizon *safety* missions, where the plant is expected to reach a goal while avoiding a majority or all of the agents, already satisfy the condition $|\mathcal{M}_1| << |\mathcal{M}|$.

VI. CONCLUSIONS

In this paper we presented a highly efficient incremental method for automatically synthesizing optimal control policies for a system comprising a plant and multiple independent agents, where the plant is expected to satisfy a high level mission specification in the presence of the agents. We considered independent agents modeled as Markov chains and assumed that the plant was modeled as a deterministic transition system. However, our approach is general enough to accommodate plants modeled as Markov Decision Processes. For mission specifications, we considered syntactically co-safe Linear Temporal Logic formulas over a set of propositions that are satisfied by the components of the system. If a probability threshold is given, our method exploits this knowledge to terminate earlier and returns a sub-optimal control policy. Otherwise, our method synthesizes an optimal control policy that maximizes the probability of satisfying the mission. Since our method does not need to run to completion, it has practical value in applications where a safe control policy must be synthesized under resource constraints. For future work, we plan to extend our approach to mission specifications expressed in full LTL as opposed to a subset of it.

REFERENCES

[1] E. A. Emerson, "Temporal and modal logic," in *Handbook of Theoretical Computer Science: Formal Models and Semantics*, J. van Leeuwen, Ed. North-Holland Pub. Co./MIT Press, 1990, vol. B, pp. 995–1072.

[2] C. Baier and J.-P. Katoen, *Principles of Model Checking*. The MIT Press, 2008.

[3] P. Tabuada and G. J. Pappas, "Linear time logic control of discrete-time linear systems," *IEEE Transactions on Automatic Control*, vol. 51, no. 12, pp. 1862–1877, 2006.

[4] M. Kloetzer and C. Belta, "Automatic deployment of distributed teams of robots from temporal logic specifications," *IEEE Transactions on Robotics*, vol. 26, no. 1, pp. 48–61, 2010.

[5] T. Wongpiromsarn, U. Topcu, and R. M. Murray, "Receding horizon control for temporal logic specifications," in *Hybrid systems: Computation and Control*, Stockholm, Sweden, 2010, pp. 101–110.

[6] J. Tumova, B. Yordanov, C. Belta, I. Cerna, and J. Barnat, "A symbolic approach to controlling piecewise affine systems," in *IEEE Conference on Decision and Control (CDC)*, Atlanta, GA, 2010.

[7] D. Berwanger, K. Chatterjee, M. De Wulf, L. Doyen, and T. A. Henzinger, "Strategy construction for parity games with imperfect information," *Inf. Comput.*, vol. 208, pp. 1206–1220, October 2010.

[8] W. Thomas, "Infinite games and verification," in *CAV*, 2002, pp. 58–64.

[9] M.Kloetzer and C. Belta, "Dealing with non-determinism in symbolic control," in *Hybrid Systems: Computation and Control: 11th International Workshop*, ser. Lecture Notes in Computer Science, M. Egerstedt and B. Mishra, Eds. Springer Berlin / Heidelberg, 2008, pp. 287–300.

[10] H. Kress-Gazit, G. Fainekos, and G. J. Pappas, "Where's waldo? sensor-based temporal logic motion planning," in *IEEE Int. Conf. on Robotics and Automation*, 2007, pp. 3116–3121.

[11] A. Bianco and L. de Alfaro, "Model checking of probabilistic and nondeterministic systems," in *Foundations of Software Technology and Theoretical Computer Science*, ser. Lecture Notes in Computer Science. Springer Berlin / Heidelberg, 1995, pp. 499–513.

[12] M. Kwiatkowska, G. Norman, and D. Parker, "Probabilistic symbolic model checking with prism: A hybrid approach," in *International Journal on Software Tools for Technology Transfer*. Springer, 2002, pp. 52–66.

[13] X. C. Ding, S. L. Smith, C. Belta, and D. Rus, "MDP optimal control under temporal logic constraints," in *2011 IEEE Conference on Decision and Control (CDC 2011)*, Orlando, FL, 2011.

[14] O. Kupferman and M. Y. Vardi, "Model checking of safety properties," *Formal Methods in System Design*, vol. 19, pp. 291–314, October 2001.

[15] S. Jha, S. Gulwani, S. Seshia, and A. Tiwari, "Synthesizing switching logic for safety and dwell-time requirements," in *International Conference on Cyber-Physical Systems*, 2010, pp. 22–31.

[16] S. Gulwani, S. Jha, A. Tiwari, and R. Venkatesan, "Synthesis of loop-free programs," in *32nd ACM SIGPLAN conference on Programming Language Design and Implementation*, 2011, pp. 62–73.

[17] A. Bhatia, L. E. Kavraki, and M. Y. Vardi, "Motion planning with hybrid dynamics and temporal goals," in *IEEE Conf. on Decision and Control*, 2010, pp. 1108–1115.

[18] L. Alfaro, *Formal verification of probabilistic systems*. Standford University, 1997, Ph.D. dissertation.

[19] T. Latvala, "Efficient model checking of safety properties," in *Model Checking Software. 10th International SPIN Workshop*. Springer, 2003, pp. 74–88.